algorithm is then used to enhance the accuracy of the speech recognizer without human supervision. This technique is useful for adapting a large vocabulary ASR engine.--

IN THE SPECIFICATION

Please replace the following paragraphs with the below written paragraphs:

The paragraph beginning on page 7, line13:

Figure 3 is a graph illustrating the results of combining two adaptation algorithms, the
HMM adaptation, and the Word Bigram Statistics adaptation.--

Page 8, in the paragraph beginning on line 13, change as follows:

The routine begins at step 20 by deploying the ASR engine, for example, at a customer installation in the field. It is assumed that the ASR engine as initially deployed is imperfect and must be tuned to increase its accuracy level. At step 22, a test is run to determine whether the ASR engine has been set for automatic adaptation according to the invention. This may be achieved, for example, by extending the engine's application programming interface (API) to include a flag that may be set (e.g., by the application developer or the user) to begin the adaptation process. The flag may be set remotely after the engine has been deployed. If the engine

has been set for automatic adaptation, the routine continues in step 24. At this step, a test is made to determine whether a given data-collection period (e.g., a 24 hour period) has elapsed. If not, the routine continues at step 26 to store the recognition results, along with the associated speech waveform samples. In particular, during this step, the recognition results obtained by processing a live input 38 are saved to a disk file. As will be seen, the recognition results may include the actual results (i.e., the hypothesized spoken utterances) generated by the ASR engine, together with information such as confidence levels, n_best hypotheses, and other data which might be used as input to the adaptation algorithms in step 30.

Page 9, in the paragraph beginning on line 9, change as follows:

In a variant of step 26, the system may be configured to save a "downstream" version of the speech data (e.g., cepstral coefficients), instead of the "raw" digitized speech waveform 39 samples. This is another advantage of the present invention. In particular, because there is no requirement for humans to listen to the speech data, significant data-reduction may be obtained by storing only the form of the speech data that is required for executing the adaptation algorithms. This advantage can result in reduction in costs for computer equipment, including CPUs, IC memory, and hard disks.

Page 10, in the paragraph beginning on line 2, change as follows:

Step 26 cycles until the result of the test at step 24 is positive. At this point, the routine continues at step 28 to retrieve the information saved during the time period. At step 30, an adaptation algorithm (or a plurality of adaptation algorithms) is executed against the information to increase the accuracy of the engine. As will be seen, this algorithm may be based on an acoustic model 33 (e.g., Hidden Markov Modeling) 34, a language model (e.g., Word Bigram Statistics) 35, a pronunciation model 36 (e.g., Phonetic Transcription) 37, or some combination of these different model types. At step 32, the so-tuned recognition engine is then re-installed in the application, presumably with better accuracy and more efficient use of computing resources than the original engine.

The paragraph beginning on page 14, line 19:

--For these experiments, the ASR engine's first-pass Viterbi search graph was biased with word bigram data extracted from subsets of the recognition results on given development input data. The result transcriptions (i.e., recognition hypotheses) were randomized and then various sized portions, starting form the to-choice down, were taken to accumulate word pair frequencies. A variation on this experiment imposed a score threshold on the recognition results as the sub-setting mechanism. These counts were converted into

probabilities, and these probabilities were used to bias the Viterbi search in favor of the most likely word sequences—

Delete the paragraph beginning on page 17, line 12:

The paragraph beginning on page 18, line 7:

--As described earlier, each of the above approaches is fairly orthogonal. Thus, two or more of these adaptation methods may be combined to produce an additive benefit. A simple combination of the HMM and Bigram adaptation generated the results illustrated in Figure 3. Combining these adaptation algorithms had a somewhat less than additive effect of the development test data, and a somewhat more than additive effect on the evaluation test data. The combined improvement in both cases approached 18%.--

IN THE CLAIMS

Please amend the claims as follows:

1. A method of improving the recognition accuracy of a speech recognizer, comprising the steps of:

deploying the speech recognizer in an environment to receive live input data;

receiving live input data ;

without supervision, applying a given adaptation algorithm to the received live input data as it is being recognized and a